REMARKS

Amendments to Specification

Two simple clerical corrections were made to the Specification.

Objections

The Office Action objected to claim 5 as being dependent on a rejected base claim, but stated that it would be allowable if written in independent form, incorporating all the limitations of the base claims and intervening claims. New independent claim 35 now incorporates the limitations of former independent claim 1 and former dependent base claim 4 into former dependent claim 5, and therefore should be allowable.

Prior Art Rejections

Comparison of Prior Art References with Present Invention

The essence of the present invention is a scalable, hierarchical approach that is capable of providing a call trace in real time. The invention operates effectively even when probes are separated by a relatively slow wide area network (e.g., the Internet) using the TCP/IP or similar protocol. We have amended the remaining pending claims to distinguish the Lin (U.S. 6,028,914), Birnahk (U.S. 6,400,813), and Swanson (U.S. 5,867,558) prior art references cited by the Examiner, both alone and in combination, and have added a number of additional claims. A general overview is in order before proceeding to a detailed discussion of the claims.

Neither Lin nor Birnahk attempts to do a call trace in a real-time. Lin goes through a complex negotiation process to select one probe to serve as an anchor. Only after anchor negotiation is completed, all probes having messages relevant to the call get their messages out of storage and send them to the anchor probe, where the call is reconstructed. Lin does not scale well. Suppose that there are 10 relevant signal messages, each detected by a different probe, for a given call, and 100,000 probes. In Lin's approach, each of the 10 probes will send out a message to every other probe, totaling nearly 1 million messages that must be sent before even a single message has

been sent back to the computer system or operator conducting the trace. And many steps in addition to the anchor negotiation need to be completed before any messages are returned to the operator: (1) each probe that has detected a relevant signaling message must wait to see if any other probe declares that the other probe detected an earlier relevant signaling message; (2) some probe must give up waiting and claim to be the anchor; (3) the anchor probe must notify all the other probes that it has claimed to be anchor (and apparently, with TCP/IP network latency, claims from competing claimant anchors could cross—a contingency that Lin does not provide for); (4) all probes must check their storage buffers (queues) to see whether they hold any relevant signaling messages; (5) each probe must send any relevant signaling messages from its respective buffer to the anchor probe; (6) the anchor probe must somehow decide it will not be getting any more relevant messages (apparently again after simply giving up waiting); (7) the anchor probe will sort and correlate the relevant messages; and, ***finally,*** send begin to send relevant messages to the computer or operator conducting the trace. All relevant signaling messages will undergone storage in probe buffers (queues) before being sent to the anchor. The timestamp sent by each probe detecting a relevant message is used by Lin solely for negotiation to choose an anchor probe.

In the approach of the present invention, the probes are all "armed" *in advance of a diagnostic telephone call* with the necessary information to detect the triggering (e.g., IAM) signaling message associated with that call. Because the probes expect the diagnostic call, signaling link messages can be detected and returned in real time. As applied to the example above, a probe of the present invention detecting the triggering signaling message will send out a message that will be ultimately be received by each of the 100,000 probes. But because of the *a priori* arming, only relevant messages need to be sent after the trigger is received. The number of messages once the call has been placed is only about 10% of Lin, or less. Once a probe has received the trigger, any new relevant messages detected in real time passing across its respective signaling links will be sent without delay and without intervening storage up to the computing device or operator conducting the trace. In order to allow connections among the component element managers and probes to be made over TCP/IP networks, which are relatively slow compared to expensive voice quality links, each probe maintains a small buffer. The timestamp sent with the trigger message broadcast to is used to determine how far back in time within its buffer a probe must look in the buffer for relevant messages. The

timestamp is not used for any kind of anchor negotiation. Each probe sends its signal link messages to its respective element manager, and each element manager knows to forward those signal link messages to the investigating computer or operator. Note that the second message (e.g., address complete message, ACM) in the sequence associated with a call is not issued in less than a fixed amount of time *D* (e.g., 100 msec) after the first one (e.g., IAM). Thus, if a probe sees that the time at which begins to watch the links for real time messages does not exceed the timestamp time in the trigger by more than *D*, then it does not need to check the buffer at all. Also note that there is a single timestamp time for a given call trace, namely the time of the triggering signaling message, and not many timestamp times as in Lin.

While Birnahk uses a hierarchical system in which monitors supervise probes, Birnahk like Lin reconstructs the call after the fact from stored messages. Combining Lin and Birnahk does not in any sense yield either the structure or the method of a real time call tracing system like the present invention.

The Amended Claims

Claim 1 was rejected as being obvious over Lin in view of Birnahk. Birnahk "teaches mediators ... operating as element managers ... for a purpose of monitoring communication links for signaling messages related to a given call." The role of the mediator in Birnahk is significantly different from that of the present invention as embodied in amended claim 1. The Birnahk probes send relevant messages to the mediator (col. 3 line 56 – col. 4 line 14). The mediator subsequently correlates the message by call, and may produce a complete call detail record. As discussed previously, neither Lin nor Birnahk does real-time call tracing. Lin collects all messages at all probes. When an anchor probe has been identified for a particular call of interest, the probes search their storage for relevant messages and forward them to the anchor probe. Birnahk collects all messages, possibly relevant to a number of different calls, at the monitor to facilitate subsequent correlation. As amended, our claim 1 has the limitation that the second probe, "upon receiving the trigger, watches respective communication links for signaling link messages that match the call parameters and transmits any matching messages, detected passing through the respective communication links, to the computing device without intervening storage in the

buffer." Both Birnahk or Lin store all the relevant messages before any are correlated, sorted, and sent to the operator. Claim 22, dependent on independent claim 1, requires that the matching messages detected by watching the links be transmitted "in real time," a related limitation not found in Lin or Birnahk.

Claim 1 is also distinguishable from Lin in how the timestamp is utilized. Lin uses the timestamp in a negotiation process to determine which probe will be the anchor to which other probes will send their relevant messages. Each probe that sees a relevant message must send a message to all other probes. Then each probe compares the time of the first relevant message it detected with relevant messages detected by other probes. Obviously, no probe can know definitively that it was the first to detect a relevant message because it never knows definitively whether it has actually received all the messages from other probes. In effect, one of them quits waiting and declares itself to be the anchor. Throughout all this, the relevant messages are just sitting idly in queues at the various probes.

The timestamp in the present invention as in claim 1 is not used to negotiate which component is in charge. The timestamp, included within the trigger, "indicat[es] the time of the initial signaling link message" (claim 1a(iii)), typically an IAM message. In other words, unlike Lin, the present invention takes advantage of the inherent meanings of the signaling link messages themselves to get the process underway and to determine the roles of the various probes and element managers. Note, however, that our approach implies that the probes must be armed *in advance* to look out for the initial signaling link message. In other words, the call trace will be made on a call that, at the start of the process, has not been yet placed. Lin and Birnahk do not have this requirement, so why is this approach useful? Suppose, for example, that a customer informs the phone company that her call did not go through. The phone company wants to determine what caused the problem. The probes are armed for a call trace, and then the operator makes the call. As specified in subsequent claims, the call trace events can be diagnosed, and problems can possibly even be repaired as they actually occur in real time.

The basic elements enabling this real-time functionality are found in amended claim 1. The probe first detecting the initial signaling link method "issues a trigger including a timestamp ... and further including the call parameters" (claim 1a(iii)). The call parameters, which the probe "extracts ... from the initial triggering signaling link

message" (e.g., IAM) could be, for example, OPC/DPC/CIC or OPC/DPC/TID combinations. The call parameters so extracted facilitate the call trace, but they cannot be known in advance of the call. Ultimately, a second probe "receives the trigger" (claim 1c(iii)). At this point, the second probe does two things—it searches its buffer for any relevant messages that might have passed through its links before the second probe received the trigger (claim 1c(iv)), and begins watching its links in real time for new messages specifically relevant to this call (claim 1c(v)); the matching messages detected on the links are transmitted "to the computing device without intervening storage in the buffer." The probe transmitting upon detection any messages observed on the links is how this invention achieves tracing in real time, a limitation on the probes not present in Lin, Birnahk (or Swanson).

While Lin uses a buffer (queue) similar to the present invention, Lin puts all messages into the buffer and correlation and sorting of relevant messages are done after the fact. In the present invention, the buffer is only used to catch those signal link messages, if any, that might have slipped through the cracks before the second probe receives the trigger. This buffer allows the present invention to achieve near real time trace even over a connection such as TCP/IP Internet, which is relatively slow compared to the expensive voice quality connections to probes used in the prior art. Again, while Lin uses a buffer, Lin does no real time transmission of relevant signaling messages without putting them first into the buffer, then negotiating for an anchor, then transmitting all messages to the anchor, then correlating the messages, and so forth.

New claims 27 - 34 should be allowed because they depend on now allowable base claim 1. These claims establish structural elements distinguishing the present invention, namely: (1) each probe is assigned to a single element manager, (2) probes need not be aware of each other, (3) element managers can communicate across a wide area network; (4) all elements can communicate over TCP/IP, and (5) a minimum amount of information is stored and transferred.

Claim 6 is an independent method claim that was rejected for lack of novelty over Lin. It has been amended analogously to the system of claim 1, and should be allowed for the same reasons. Claims 36 - 49 should be allowed because they depend on base claim 6, and for the reasons stated above in connection with claims 21 - 34.

The rejections of claims that have been canceled are moot.

CONCLUSION

All of the claims remaining in this application should now be seen to be in condition for allowance. The prompt issuance of a notice to that effect is solicited.

Respectfully submitted,
COMPUTER NETWORK
TECHNOLOGY CORPORATION
By its attorneys:

Date: 11/1/2005

James C. Evans
Registration No. 56,730
Beck & Tysver, P.L.L.C.
2900 Thomas Ave., #100
Minneapolis, MN 55416
Telephone: (612) 915-7006
Fax: (612) 915-9637